Structuring Peer-2-Peer Communities

Asif Akram and Omer F. Rana
Department of Computer Science
Cardiff University, UK
{A.Akram, O.F.Rana}@cs.cf.ac.uk

Abstract

Locating suitable resources within a Peer-2-Peer (P2P) system is a computationally intensive process, with no guarantee of quality and suitability of the discovered resources. An alternative approach is to categorise peers based on the services they provide – leading to the interaction between peers with common goals to form societies/communities. Organization of peers in different communities is suggested to be useful for efficient resource discovery. We analyse the types of communities that may be useful, and how they may be structured.

1. Introduction

Emerging distributed computing paradigms, such as Grid Computing, comprise of dynamic and distributed resources/peers organised as a "Virtual Organisations". Resource discovery plays a significant role in organising and managing such organisations. Resource discovery becomes a time-consuming process and imposes an overhead on network access. The numbers of interactions are likely to increase exponentially as the numbers of peers grow. Restricting interactions to be between a set of peers is a key factor to scale the resource discovery problem. Peers can be categorised based on criteria such as the type and quality of service they provide, etc. Any initial cost in categorising peers can provide benefits for discovering *preferable* peers without a large discovery cost subsequently – thereby leading to the development of "communities".

The concept of communities is similar to interactions between different departments at a University. For instance, a lecturer can be a member of different faculties e.g. a mathematics lecturer teaching calculus to computer science students. A similar problem in Grid Computing is what Davis and Smith refer to as the "connection problem" [1], where peers need to find other suitable peers to co-operate with, assist, or interact with. "Focused Addressing" [2] is one solution to the connection problem where requests are sent to particular subset of peers, believed to assist the requesting peer. We provide an alternative solution based on structuring communities.

Individual peers, although selfish, are expected to interact with each other in some way. Co-operation of one form or another therefore becomes essential. Each peer prefers to be in an environment where it may be easily discovered by a suitable user, and can locate other peers with minimum efforts. Peers providing different services may be grouped together based on attributes such as type of services, resources owned and domains of operation. Each community has one Service Peer with dual responsibility of not only managing the member peers but also keeping track of other communities with which they interact on behalf of member peers. Interaction between communities is only through the Service Peers, except in "ad-hoc" communities. Direct interactions between peers restrict message traffic to enable scaling within a Grid system.

2. Community Formation

When a new peer joins the network, it tries to discover the Service Peer which may have interest in the capabilities/services provided by the new peer. If the interests of a Service Peer are different, the new peer is referred to other Service Peer/s, or the new peer tries to locate alternative Service Peer/s with compatible interests. A Service Peer and all peers registered with it constitute a community. A Service Peer manages all peers within the community and communicates with neighbouring Service Peers on the behalf of member peers. A Service Peer is essential for the bootstrapping of a new peer, as it supports a new peer to discover enough network resources to sustain itself. A Service Peer may interact with a monitoring service within a community to achieve this. We therefore also foresee the existence of common infrastructure services (such as monitoring, directory, security/certificate authority, etc) within each community.

3. Type of Communities

Individual autonomous peers have expertise and interests in specific resource/s. Based on these expertise and interests, peers are grouped together, but expertise and interests are not the only criteria for categorizing peers. Communities/societies can be of different types as mentioned below:

Competing Community: In a Competing Community each peer has the same expertise – although some service attributes may vary. Similarity in services may develop

competition amongst member peers, as member peers will compete each other to get selected by a client.

Co-Operative Community: In Co-Operative communities all peers provide different services, which must be used alongside services of other member peers. In such communities, each peer is dependent on at least one other member peer. Hence, when one peer is selected, then the possibility of selecting another member peer providing utility service/s is increased. This mutual co-operation is suitable for peers which provide simple services.

Goal Oriented Community: Here a collection of peers work together to achieve a particular goal. Membership in such a community is only allowed to accomplish the assigned task. Goal oriented communities may also be important in self-organising systems, where interactions between member peers are not pre-defined, but the services required are. In such instances, member peers may interact with each other in arbitrary ways to achieve a given end result.

Ad Hoc Community: Here peers can be in a co-operative or competing community, but need to work together as a team. In ad hoc communities peers interact directly with each other without interference and involvement of a Service Peer. Peers belonging to different communities providing different but supporting services form the basis of an ad hoc community, as long as both concerned communities have agreed to use each other's service.

Domain-Oriented Community: Such a community is formed by linking together similar-minded organisations and institutions, instead of the services they provide, such as academic communities, research communities, and open-source communities. Hence these communities are domain-oriented rather than service-oriented.

4. A Prototype System

We have implemented a simulation of such communities using JXTA. In the JXTA prototype there is an option for creating Groups and Peers along with their descriptions, to influence membership of a community. When a Peer applies for membership, its description is matched with the description of the Group.

Each JXTA Group has a sorted list of member Peers and each Peer has a sorted list of Groups to which it belongs. Peers apply for membership based on description of the Group. At any time any Group can have five members and any Peer can be member of three different Groups (in the prototype). Peers can be added to different Groups using an interface, but membership will be awarded based on the selection criteria of that specific Group. Similarly any Peer can resign from any Group at any time.

Each Peer has its own thread, and after certain time interval it discovers new Groups from the local cache of the JXTA environment and applies for the membership. Peers monitor new Groups, and on discovering a more suitable Group, may leave their existing Group. Similarly each Group prefers to have Peers with the most compatible description, and on the membership of any new Peer may cancel the membership of existing Peer. Result of this simulation was quite encouraging and as expected in the beginning the system has Groups and Peers attached without any uniform pattern, but over time the system achieves stability, and forms Groups with compatible Peers. Once the whole system is stable, creating new Groups or Peers does not affect the overall memberships of Groups and Peers. In the prototype there is an option to change the description of Peers and Groups, which may de-stabilize the system. However, as Groups have their own thread and monitor their description with member Peers, this de-stabilization is temporary and system tends to achieve its stable state over a short time frame.

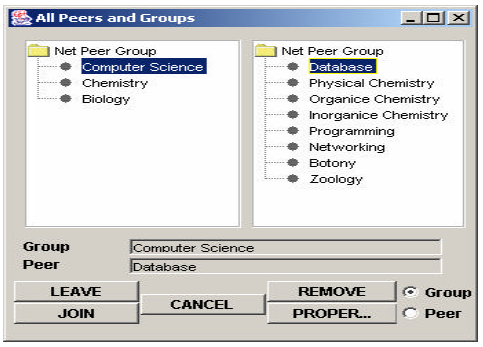

Figure: Interface to "Community Builder" in JXTA

5. Conclusion and Summary

Categorization of peers in communities on the basis of their expertise and interests is presented. Organizing peers in one form or another makes the discovery of resources efficient, with less resource consumption. Categorizing the peers in communities is simple, open and easy to implement, and the initial overhead of developing communities pays off at the time of resource discovery. Communities are more stable, and stability increases with passage of time, and are more adaptive to operate in a dynamic environment.. Communities can be of different types representing different types of human social networks. We believe that organizing peers into such communities is a useful undertaking to support problem solving in emerging distributed environments, such a Computational Grids.

References

[1] Davis, R. and R. G. Smith: 'Negotiation as a Metaphor fro Distributed Problem Solving'. Artificial Intelligence 20, 63-109, 1983.

[2] Parunak, H. V. D.: 'Distributed Artificial Intelligence', Chapt. Manufacturing Experionce With the Contract Net, pp. 285-310, Research Notes in Artificial Intelligence. Los Altos, CA: Morgan Kaufmann Publishers, 1987.